Inadvertently Programmed Bits in Samsung 128Mbit Flash Devices: A Flaky Investigation

Gary Swift
Jet Propulsion Laboratory / California Institute of Technology,
Pasadena, CA
gary.m.swift@jpl.nasa.gov

Abstract- JPL's X2000 avionics design pioneers new territory by specifying a non-volatile memory (NVM) board based on flash memories. The Samsung 128Mb device chosen was found to demonstrate bit errors (mostly program disturbs) and block-erase failures that increase with cycling. Low temperature, certain pseudo-random patterns, and, probably, higher bias increase the observable bit errors. An experiment was conducted to determine the wearout dependence of the bit errors to 100k cycles at cold temperature using flight-lot devices (some pre-irradiated). The results show an exponential growth rate, a wide part-to-part variation, and some annealing behavior.

TABLE OF CONTENTS

1. BACKGROUND

For obvious reasons, spacecraft designers would prefer to use known high-reliability military- or aerospace-grade electronic components. However, the rich, advanced features of "cutting edge" commercial IC's represent an almost irresistible siren call to the designer. An interesting case in point is memory devices. Prior to 1991, commercial DRAM's were forbidden (almost unthinkable) for in-space use. Note that the term "commercial DRAM" is redundant for all practical considerations: to this day, no DRAM manufacturer targets the military or high-rel market. However, the attractions of density (~four times commercial SRAM's and ~16 times hardened SRAM's) have proven irresistible: all JPL projects designed since then have used arrays of DRAM's either for mass storage (replacing magnetic tape) or for computer main memory (replacing SRAM's) or both.

The characteristic –density– that drove adoption of commercial DRAMs for space applications now favors flash memory. In addition, non-volatility offers irresistible power-management advantages, particularly for recorder-type applications. There are some significant disadvantages, including slow access speed (especially for erasing and writing), large block granularity for erasing, lifetime cycle limits, and for most types, some initial bad blocks. However, none of these are insurmountable and the ubiquitous use of flash memory in space seems inevitable in the absence of a "showstopper."

One potential showstopper is radiation; the radiation environment in space is considerably more severe than it is on the ground because of the shielding effects of earth's atmosphere and magnetic field. Potential radiation problems tend to fall into three categories: (1) performance degradation or failure due to total dose accumulation, (2) high-current, potentially destructive latchup events from a single ionizing particle triggering a SCR structure inherent in CMOS, and (3) particle-induced bit flips known as single-event upset (SEU). The total dose and latchup problem have been overcome so far for DRAM's through the fortuitous radiation tolerance that can often be found by surveying the large number of manufacturers that have been available. SEU problems are overcome

The research done in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration (NASA), under the NASA Electronic Parts and Packaging Program (NEPP), Code AE.

with system-level error detection and correction (EDAC).

Examples of JPL's successful application of this approach for commercial DRAM's include both data recorder and main memory designs. The *Cassini* spacecraft currently on its way to Saturn's moon Titan uses a pair of DRAM-based data recorders built by TRW. Each recorder has an array of three hundred twenty 4Mb OKI DRAM's with controllers incorporating an EDAC word consisting of 32 data bits and a 7-bit Hamming code for correcting single SEUs and detecting doubles. A very similar recorder design by SEAKR Engineering based on 4Mb Hitachi DRAM's is currently in orbit around Mars in the Lockheed-Martin built spacecraft *Mars Global Surveyor (MGS)*. Many JPL missions, from 1996's successful rover-landing *Mars Pathfinder* through current missions *MGS* and *Genesis* and next year's *Mar Exploration Rover (MER)* pair, depend on RAD6000-based flight computers from BAE Systems which use forty 16Mb IBM DRAM's for their main memory.

Mapping this DRAM experience to flash memory seems to be a logical step. Indeed, as a part of JPL's development of the next generation of avionics under the X2000 program, a non-volatile memory (NVM) board specification was drafted in 1998 assuming most (or all) of the storage would be flash. Recently SEAKR delivered flight-ready 2Gb NVM boards each populated with twenty Samsung 128Mb flash devices and a Honeywell ASIC controller using an 80-bit Reed-Solomon-based EDAC word with 64 data bits. The EDAC is capable of correcting byte-confined errors up to eight bits long and detecting errors in two bytes. This EDAC was incorporated to fix SEU's, but is powerful enough to allow full functionality (albeit without SEU correction) in the event of failure of an entire flash device. *Deep Impact*, a comet-probing mission to be launched in 2003, will be the first JPL spacecraft to fly the X2000 NVM board, and thus, the first to depend on commercial flash memory for a mission-critical application.

The response of single-voltage flash memory devices to the space radiation environment has proven to be very problematic, if not quite a showstopper. While the EDAC handles SEU's and the selected Samsung flash was determined to be immune to single-event latchup, the rapid accumulation of total dose damage from ionizing radiation is a major problem for all tested flash devices, including Samsung's. The problem stems from the on-chip charge pump used to obtain the high voltages needed for erasing and programming. The typical charge pump will not supply very much current, but leakages caused by chip-wide total dose increase the amount of current they must supply in order to maintain voltage. In combination with direct damage to the charge pump, this leakage causes a biased Samsung flash to fail to program (that is, act as a read-only memory) after only a few kilorads.

These total dose problems can be solved for some missions and radiation environments with a combination of power management and radiation shielding. The former takes advantage of the flash's non-volatility and the observed reduction in total dose damage for unbiased devices. For the selected Samsung devices, the reduction in damage is roughly a factor of three so that the addition of power-removal circuitry to the board allows operational control over the duty cycle to manage the damage. Radiation shielding can be effective for some environments, like the magnetically trapped electrons in the vicinity of Jupiter's interesting moon Europa. The NVM board accommodates a specially designed tungsten clamshell around the flash array, but note that this adds significantly to the board mass obviating one of the flash density advantages.

Two unexpected, non-radiation potential showstoppers were discovered by SEAKR when exercising a development board: "spontaneous" bit errors and block-erase failures. Both problems were increasing with additional cycling, raising again questions about the viability of using flash for space applications. The rest of this paper reviews the subsequent investigations and their conclusions. Although specifically directed at answering the engineering question of "how bad can these problems get?", the investigations do reflect on the physics question of "what is(are) the source(s) of the problems?"

2. FLAKY BITS AND BEFs

This section discusses initial, exploratory investigations. These investigations were aimed at determining the main parameters affecting error occurrence and the direction of those dependences. Before presenting those results, it is helpful to briefly review particulars of 1) the

target flash device, 2) the test board architecture, and 3) the error terminology used here. The chosen Samsung 128Mb flash is the KM29U128T (re-named in late 1999 as K9F2808UOM-YCBO). This is a 3.3V, one million cycle flash memory, organized as 1024 erasable blocks of 32 pages of 528 bytes of NAND-accessed floating-gate transistors. The test board has two independent banks of ten flash devices, designated Side A and Side B. All ten parts on a side are accessed simultaneously and together form the eighty-bit EDAC word. The controlling FPGA is programmed to run each side at near the flash's maximum speed of 2-3 ms for a block erase, 200-500 us for a page program, and 50 ns for sequential byte reading after 10 us for a page buffer load. The result is that a full erase-write-read cycle can be accomplished on a side in about 30-31 seconds. The writes are done in two half-page steps.

It was observed very quickly that almost all the bit errors were in cells that were not intended to be programmed and that some of these cells read out as programmed in a large fraction of the cycles in which they were intended to be left erased. After confirming that cells were erasing, these errors were dubbed "inadvertently programmed," and since any cell that had previously shown an error was more likely to be in error again, these are called "flaky bits." On the other hand, block-erase failures or BEF's are indicated by a status bit read immediately after an erasure. Although Samsung specifically admonishes against using these ("If erase operation results in an error, map out the failing block..."), in practice, they typically read out as fully erased and appear to program normally. Thus, the state machine appears to apply a tougher criterion for fully removing charge from a cell's floating gate than the read-out sense amps. Because blocks that exhibit a BEF remain quite usable (and, incidentally, because the flight-board design discards the erase status), BEF's are considered less serious than bit errors. It seems a practical solution to mapping out "bad" blocks to apply a bit error criterion, such as designating a block bad when one or more flaky bits consistently are erroneous.

The parameters crudely explored in the SEAKR preliminary investigations were: temperature, bias, cycling, data pattern, and wafer lot. Also, independent verifications of the error modes were made at JPL and at Semiconductor Solutions on individual devices (without the use of the test board). These results, which involved only a few thousand cycles, can be summarized as follows:

1) Colder is worse. Note that the cold temperature used was –55 degrees C, well below the part spec. limit.
2) Lower bias is worse.
3) Both the number of flaky bits and the likelihood a flaky is in error increase with cycling.
4) Pseudo-random patterns were worse than checkerboards.
5) Of the two lots used, the Side A lot was significantly worse (see Fig. 1A & B).

Two factors likely confuse the interpretation. First, as stated above, it is clear from Fig. 1 that the dominant physical mechanism causing bit errors is different in the two lots of devices used in the exploratory investigations. Second, none of these investigations make a distinction between factors affecting *damage accumulation* and those affecting the *observation of the damage*. It is likely that some parameters may affect the two in opposite directions, but the investigations only observe the *net* effect. For example, it seems likely that higher bias will cause more wearout, but the observation of more errors at lowered bias may be the result of shifting the sense amp's detection threshold. The observed temperature dependence could also be an example. The observation of more errors at lower temperature could be explained by either increased inadvertent charging or from shifting minimum detection level against a fixed distribution. (It could also be the result of tilting the competition between wearout and annealing towards the former).

It can be inferred from spec. sheet changes made in April 1999 [4} and an app. note from late 1999 [5]that Samsung was aware of and wrestling with bit error problems, too. The spec. sheet [4] now included a technical note on "ECC (error correcting code)" needed... "to implement a highly reliable system." The application note [5] has sections on "ECC Design Guide" and "Invalid Block(s) Management." Perhaps most interesting is the change in the maximum number of writes to a page from 10 to 2. One may infer that this is to minimize program disturbs that were occurring more frequently than previously expected.

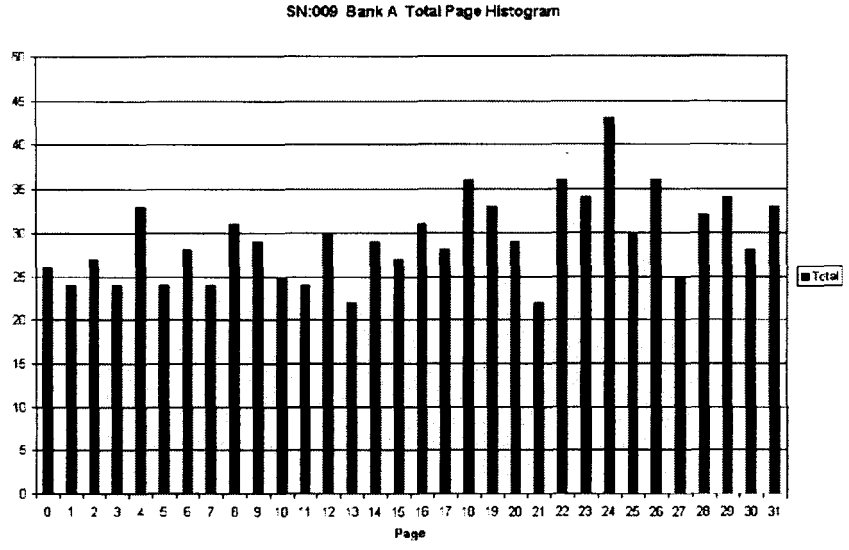

Fig. 1a: Histogram of Bank A for serial number 9

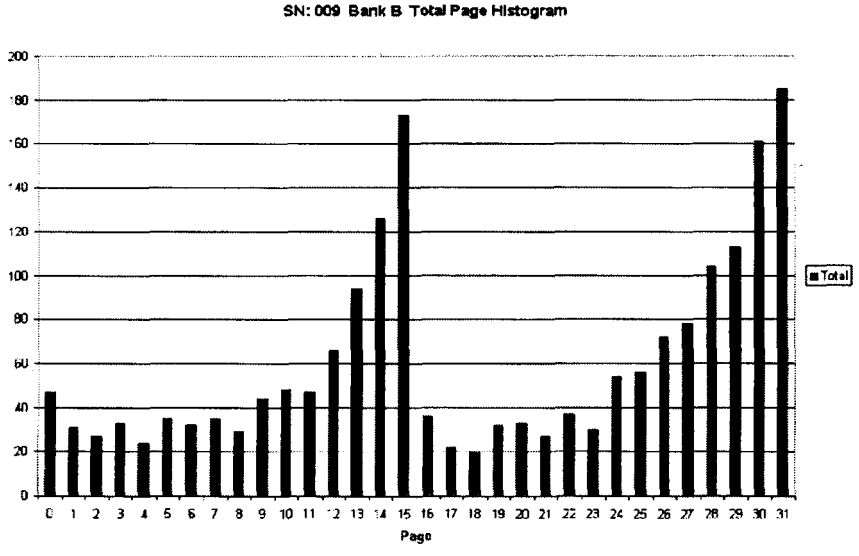

Fig. 1b: Histogram of Bank B for serial number 9.

3. THE CYCLING EXPERIMENT

This section describes the results of a test intended to measure the number of flaky bits as a function of erase-write-read cycles for a representative set of flight-lot parts under worst-case temperature conditions of –55° C (below the spec. minimum of 0° C). Both sides of the FPGA-based development board were populated, each with ten flight-lot parts (marked with date code 934)– eight of the devices on Side A were previously irradiated with 3.5 krad(Si) of Co60 gammas under dynamic bias (about half the dose to functional failure due to charge pump

degradation). The experiment terminated when 100,010 cycles (one tenth of the part spec maximum) were reached. The experiment which required 833 hours of cycling time (about five weeks) was conducted over a four month period; during idle times, the devices were unbiased and ambient temperature. Figures 2a and 2b show the accumulation of bit errors and block-erase fails for Side A and B, respectively, over the entire experiment as a function of cycle.

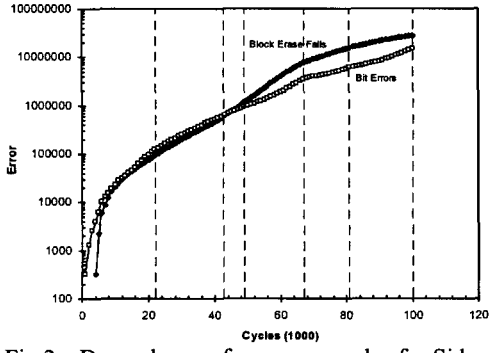

Fig 2a: Dependency of errors on cycles for Side A

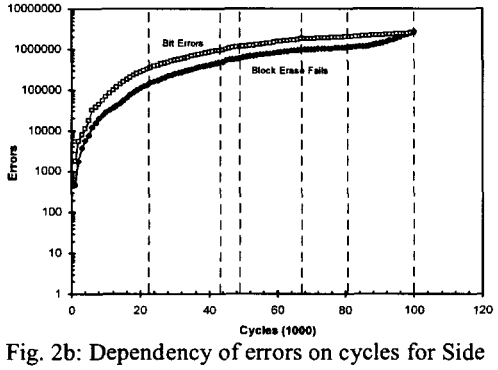

Fig. 2b: Dependency of errors on cycles for Side B

During full speed cycling, any words with errors as detected by the EDAC on read were counted. At six irregular points in the cycling, a slower set of cycles were run to record the particular block numbers with erase fails and also bit error locations and values. These characterization stops are indicated on Figure 2. Note that the number of blocks mapped out is sometimes increased at these spots, as indicated on Figure 2. Through cycle 49080, only a composite of the original manufacturer's bad block lists were mapped out – 13 on Side A and 20 on Side B. After that the number mapped out increased significantly going to 366 on Side A and 91 on Side B. Note the number mapped out is higher during the characterization cycles, likely biasing the results down some in order to more faithfully simulate in-space use.

Table 1 presents the same results in terms of the average occurrence rate (per cycle) at each characterization stop for both sides. Note that the adjustment from the raw bit errors in Table 1b for the fraction of blocks mapped out is not conservative (but better than no adjustment) because the blocks mapped out probably have more errors than the remainder which were counted (that was the map out criteria).

The extra visibility of characterization cycling allows errors observed to be broken down by chip number (see Table 1 for chip numbering scheme). These details are presented for Side A at each Stop in Table 2a and for Side B in Table 2b. Note that there is a wide variance part-to-part.

The number of flaky bits observed for each chip is given in Table 3a for Side A and Table 3b for Side B. There is some inconsistency in the number of flaky bits, but in general the number of flaky bits increases with the number of cycles. This leads to the conclusion that, with increasing wearout, the flaky bits continue to incur damage that causes them to produce errors more often.

Finally, Figure 3 shows how the number of flaky bits increases with the number of E/W cycles. The number of errors was adjusted for visibility and pattern. Note the much larger error rate for the abnormal device (unit 9) compared to typical devices.

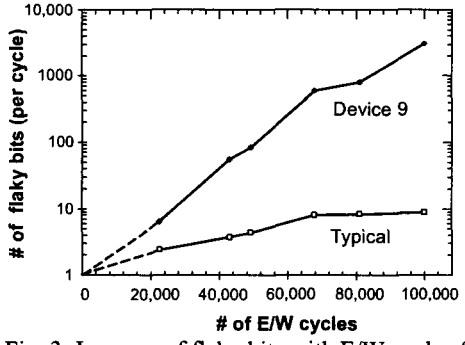

Fig. 3: Increase of flaky bits with E/W cycles for device 9 and a typical device

Table 1a: Average Block Erase Fails per Cycle by Side		
at Cycle #	Side A	Side B
22480	9.9 OF 1011	12.3 of 1004
43080	62.8 of 1011	21.3 of 1004
49080	130.1 of 1011	22.8 of 1004
67730	593 of 999	21.1 of 999
81010	724 of 808	15.2 of 952
100010	657 of 658	198 of 933

Table 1b: Average Bit Errors per Cycle by Side				
	Side A		Side B	
at Cycle #	raw	BVF adj	raw	BVF adj
22480	12.7	12.9	24.6	25.1
43080	42.1	42.6	35.7	36.4
49080	60.9	61.7	37.6	38.3
67730	325.9	334.1	41.6	42.6
81010	348.8	442	26.9	28.9
100010	1026.7	1598	32.8	36

"BVF adj" means adjusted for block visibility factor (due to mapped out blocks)

Table 2a: Number of Errors by Chip – Side A (irradiated)											
	Chip #										
at Cycle #	0	1	2	3	4	5	6	7	8	9	total
22480	9	0	6	3	n/v	9	2	1	13	n/v	43
43080	9	3	0	3	n/v	4	1	6	24	n/v	50
49080	85	11	19	10	2	22	23	20	145	515	852
67730	61	10	44	0	5	10	61	40	108	2655	2994
81010	43	10	22	8	0	17	20	62	80	3100	3362
100010	60	0	0	21	11	15	0	49	2	6369	6527

n/v means no visibility

Table 2b: Number of Errors by Chip – Side B (unirradiated)											
	Chip #										
at Cycle #	0	1	2	3	4	5	6	7	8	9	total
22480	8	7	9	19	n/v	2	1	7	3	n/v	56
43080	24	2	30	27	n/v	6	12	10	1	n/v	112
49080	133	66	94	79	32	47	21	29	7	6	514
67730	202	64	156	43	27	89	1	73	25	25	705
81010	104	3	166	65	54	31	10	41	6	57	537
100010	54	32	78	85	51	66	25	38	10	36	475

n/v means no visibility

Table 3a: Number of Flaky Bits by Chip – Side A (irradiated)											
	Chip #										
at Cycle #	0	1	2	3	4	5	6	7	8	9	total
22480	5	0	2	2	n/v	5	1	1	4	n/v	20
43080	2	2	0	2	n/v	2	1	2	6	n/v	17
49080	8	3	4	3	2	7	4	3	15	105	154
67730	4	1	3	0	2	5	5	2	6	229	257
81010	3	3	3	4	0	3	1	4	4	271	296
100010	3	0	0	2	3	4	0	3	1	625	641

n/v means no visibility

Table 3b: Number of Flaky Bits by Chip – Side B (unirradiated)											
	Chip #										
at Cycle #	0	1	2	3	4	5	6	7	8	9	total
22480	4	3	6	12	n/v	2	1	4	3	n/v	35
43080	11	2	14	11	n/v	3	6	3	1	n/v	51
49080	27	10	29	23	12	8	10	6	6	5	136
67730	18	4	16	10	9	9	1	6	5	7	85
81010	11	1	21	12	12	6	3	5	3	9	83
100010	13	4	16	13	10	8	5	5	2	9	85

n/v means no visibility

4. STEPS TOWARD A MODEL

In the absence of a physical model for the bit and block errors observed, the results of this test are more useful when interpreted in light of an empirical model which is consistent with the data from this test and also incorporates reasonable assumptions derived from other tests and the literature. This section documents the current empirical model, but note that it is somewhat speculative.

A flaky bit is damaged to some extent beyond a threshold condition within the circuit so that it sometimes shows an error. Although there is some evidence that the damage may anneal, this process is slow enough to neglect. The damage accumulates with each erase cycle, but not with reading. Total dose damage (added leakage and transistor threshold shifts) is essentially different from flaky damage, although device changes from total dose can raise the threshold for observing an error due to flaky damage . An error appears when an unprogrammed flaky bit reads back as programmed. When accomplice bit(s) are programmed, some charge also accumulates on the connected flaky bit's floating gate. The amount of charge deposition from this process is somewhat probabilistic. A flaky bit that is supposed to be unprogrammed (one) has a non-zero probability of being charged enough from the intentional programming of accomplice bit(s) to readback as zero. Thus, a flaky bit has four possible states: (1) intentionally programmed, i.e., correctly storing a zero, (2) unprogrammed with accomplice bit(s) also unprogrammed, i.e., correctly storing a one, (3) unprogrammed with small accomplice charging, i.e. deterministically reads out correctly as a one, or (4) unprogrammed with a large enough accomplice charging to deterministically read out erroneously as a zero. In rare cases, a flaky bit may be inadvertently programmed to the setpoint of the sense amp and has a non-zero probability for both a zero and one readback.

Block-erase fails occur because some floating gates are not fully discharged by the erasure process. The fraction of these not-fully-discharge bits that are also flaky bits are more likely to be charged to above the erroneous readback threshold by programming accomplice bits. Thus, block-erase fail blocks will tend to show more flaky bits in error more often.

Low temperature mainly lowers the threshold charge a floating gate needs to be read as programmed. Thus, flaky bits that are not charged enough to read out erroneously at high temperature will show errors at low temperature.

Note that overall the model is supported (or not contradicted) by all the data analyzed, but other models may also agree with the somewhat limited data.

5. CONCLUSIONS

In conclusion, the engineering test of flight lot devices to 100,000 erase/write cycles was very successful and indicates that the flaky bit problem is not likely to have a severe impact on in-space performance. A significant part-to-part variation was observed in that one device (of 20 tested) was many standard deviations worse than the others. The results indicate that radiation does not exacerbate the bit error rate, although it may contribute to more rapidly increasing the occurrence of block-erase fails. The test enabled a semi-empirical model to be proposed. However, it is important to note that 1) there are

likely other models that could explain the same data, 2) it is not known how specific the results are to the 1998-vintage flight-lot devices, and 3) even if the model is correct as far as it goes, it does not identify the root cause or a solution to the problem of inadvertently programmed bits.

6. ACKNOWLEDGEMENT

Special thanks go to Sam Day of SEAKR for executing the testing and especially for his attention to detail, and to Steven Guertin of JPL for his help in corralling and dissecting the data and for many fruitful discussions.

REFERENCES

D. Nguyen et al., "Radiation Effects on Advanced Flash Memories, *IEEE Transactions on Nuclear Science*, vol. 46, no. 6, p. 1744 (Dec.1999)

D. Nguyen et al., "Total Dose Effects on Flash Memories," *Proceedings of the 1998 Nuclear and Space Radiation Effects Data Workshop*, IEEE Pub. No. 98TH8385, p. 100 (July 1998)

G Swift et al., "In-Flight Observations of Multiple-Bit Upset in DRAMs, *IEEE Transactions on Nuclear Science*, vol. 47, no. 6, p. 1744 (Dec. 2000)

Samsung Electronics, Memory Product and Technology Division, "Application Note for NAND Flash Memory (Revision 2.0)," Dec. 12, 1999.

Samsung Electronics, Specification, "16M x 8 Bit NAND Flash Memory: K9F2808U0M-YCB0, K9F2808U0M-YIB0," Rev 1.0 July 14, 1998 and Rev. 1.1, April 10, 1999. (Now at Rev. 1.4, July 17, 2000)

AUTHOR BIO

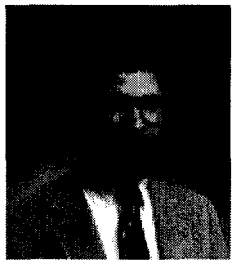

Gary Swift received his B.S. in engineering physics from the University of Oklahoma in 1975 and did his graduate work in nuclear engineering at the University of Illinois at Urbana-Champaign.

He joined the Jet Propulsion Laboratory in 1991 and is a Senior Member of the Radiation Effects Group of the Electronic Parts Engineering Office.

He has published more than 20 refereed papers, including work on dielectric rupture in commercial DRAMs and field-programmable gate arrays used on spacecraft. He has a broad range of single event, total dose and dosimetry interests with particular emphasis on efficient radiation screening of commercial microelectronics for space applications, microdosimeters, mitigation techniques for radiation effects in space, and understanding the physical mechanisms of radiation effects in emerging microelectronics. Gary and other group members received the JPL/NASA Excellence Award in 1998. They also received the Outstanding Paper Award at the 1999 Nuclear and Space Radiation Effects Conference for their paper entitled "Angular and Energy Dependence of Proton Upset in Optocouplers" which explained how direct ionization from protons increases the upset rate in optocouplers.